AFRL-IF-RS-TR-2001-117
Final Technical Report
June 2001

SECURE MULTICAST PROTOCOLS FOR GROUP COMMUNICATION

University of California, Santa Barbara

Sponsored by
Defense Advanced Research Projects Agency
DARPA Order No. C929

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK

20010810 033

This report has been reviewed by the Air Force Research Laboratory, Information Directorate, Public Affairs Office (IFOIPA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

AFRL-IF-RS-TR-2001-117 has been reviewed and is approved for publication.

APPROVED: CHESTER J. MACIAG
Project Engineer

FOR THE DIRECTOR: WARREN H. DEBANY, JR., Technical Advisor
Information Grid Division
Information Directorate

If your address has changed or if you wish to be removed from the Air Force Research Laboratory Rome Research Site mailing list, or if the addressee is no longer employed by your organization, please notify AFRL/IFGB, 525 Brooks Rd, Rome, NY 13441-4505. This will assist us in maintaining a current mailing list.

Do not return copies of this report unless contractual obligations or notices on a specific document require that it be returned.

SECURE MULTICAST PROTOCOLS FOR GROUP COMMUNICATION

Louise E. Moser
P.M. Melliar-Smith

Contractor: University of California
Contract Number: F30602-95-1-0048
Effective Date of Contract: 24 July 1995
Contract Expiration Date: 24 August 1998

Short Title of Work: Secure Multicast Protocols for Group Communication
Period of Work Covered: Jul 95 – Aug 98

Principal Investigator: Louise E. Moser and P.M. Melliar-Smith
Phone: (805) 893-4897
AFRL Project Engineer: Chester Maciag
Phone: (315) 330-3184

Approved for public release; distribution unlimited.

This research was supported by the Defense Advanced Research Projects Agency of the Department of Defense and was monitored by Chester Maciag, AFRL/IFGB, 525 Brooks Rd, Rome, NY.

REPORT DOCUMENTATION PAGE	Form Approved OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY *(Leave blank)*	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	Jun 01	Final Jul 95 - Aug 98

4. TITLE AND SUBTITLE

SECURE MULTICAST PROTOCOLS FOR GROUP COMMUNICATION

5. FUNDING NUMBERS

C - F 30602-95-1-0048
PE - 62301E
PR - C929
TA - 01
WU - 03

6. AUTHOR(S)

Louise E. Moser and P.M. Melliar-Smith

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Dept of Electrical and Computer Engineering
University of California, Santa Barbara 93106

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Defense Advanced Research Projects Agency
3701 North Fairfax Drive
Arlington VA 22203-1714

AFRL/IFGB
525 Brooks Rd
Rome NY 13441-4505

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

AFRL-IF-RS-TR-2001-117

11. SUPPLEMENTARY NOTES

AFRL Project Engineer: Chester Maciag, IFGB, 315-330-3184

12a. DISTRIBUTION AVAILABILITY STATEMENT

Approved for public release; distribution unlimited.

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)*

This final report for this project focuses on: Design and implementation of the Secure Group message ordering and group membership protocols, which are derived from our previous Trans/Total protocols. Design and implementation of the Secure Ring message ordering and group membership protocols, which are derived from our previous Totem single-ring protocols. Experimentation with the FORTEZZA card, the Cryptoki interface, and the Cryptolib software. Development of a system, called the Immune System, that combines the Secure Ring Protocols with the replication manager and majority voting algorithms of Eternal System (also funded by DARPA). Understanding the problem of achieving consensus (which underlies the message ordering and group membership algorithms of group communication systems) for environments that are subject to Byzantine (arbitrary) faults and malicious attacks.

14. SUBJECT TERMS

Secure Multicast Protocols for Group Communication

15. NUMBER OF PAGES

28

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UL

Secure Multicast Protocols for Group Communication

1 Significant Accomplishments and Progress

This is the final report for this project, which has focused on:

- Design and implementation of the Secure Group message ordering and group membership protocols, which are derived from our previous Trans/Total protocols.
- Design and implementation of the SecureRing message ordering and group membership protocols, which are derived from our previous Totem single-ring protocols.
- Experimentation with the FORTEZZA card, the Cryptoki interface, and the Cryptolib software.
- Development of a system, called the Immune System, that combines the SecureRing Protocols with the replication manager and majority voting algorithms of the Eternal System (also funded by DARPA).
- Understanding the problem of achieving consensus (which underlies the message ordering and group memberhip algorithms of group communication systems) for environments that are subject to Byzantine (arbitrary) faults and malicious attacks.

1.1 Design of Secure Multicast Protocols

The principal investigators of this proposal have spent many years in developing multicast group communication protocols (see Papers 11 and 13 for an overview). This project has taken that work much farther in developing secure multicast protocols that can withstand Byzantine (arbitrary) faults and malicious attacks.

Much of the computational cost of a secure multicast protocol arises from the need to compute a digital signature for each message sent and for each message received. Another cost is the need to include the signature for a message in every acknowledgment for that message. This increases the size of the acknowledgments and the cost of manipulating them. Our approach to the design of secure multicast protocols is to reduce the number of extra messages required by the protocol for acknowledgments, authentication, etc.

1.2 The Trans/Total Protocols

The Trans/Total protocols provide reliable totally ordered delivery of messages within a broadcast domain. The Trans protocol piggybacks acknowledgments on the backs of messages to reduce the number of messages broadcast. From these acknowledgments, a partial order graph is constructed. The Total protocol converts this partial order into a total order using a multi-stage voting algorithm, without any additional messages broadcast. Four membership protocols have been developed that operate on top of the Total protocol.

The Trans/Total protocols are unique in that they continue to deliver messages in a total order despite failed processors. When failures occur, other protocols stop delivering messages until a membership protocol has detected the failed processor and has reconfigured the system to exclude it. In contrast, the Trans/Total protocols continue to transmit, order, and deliver messages without a hiatus even under malicious attacks.

The Trans/Total protocols have the lowest communication cost of any totally ordered multicast protocols known to us; however, they are quite computationally expensive.

1.3 The Secure Group Protocols

The Secure Group Protocols are derived from the Trans/Total message ordering and membership protocols. The acknowledgment and voting mechanisms of the protocols, used in conjunction with digital signatures, prevent a malicious processor from disrupting the totally ordered delivery of messages, and from disrupting the membership protocol.

The Secure Group Protocols employ a public key cryptosystem in which each processor possesses a private key known only to itself with which it can digitally sign messages. Each processor also possesses the public keys of other processors with which it verifies signed messages. With high probability, the digital signatures prevent a Byzantine (malicious) processor from originating a message purporting its source to be some other processor. To reduce the cost of signing a message, only a digest of the message is signed. If a Byzantine processor sends two different messages to different destinations, purporting that they are the same message (*i.e.*, have the same message identifier) then, with high probability, the digests of the messages are different and a destination can recognize the messages as distinct and can process them as distinct messages. The messages may also be encrypted, but that encryption is orthogonal to the protocols.

The Secure Group Protocols comprise:

- Secure Trans Protocol – a reliable broadcast protocol which ensures, with high probability, that every message received by any non-faulty processor is received by every non-faulty processor. Acknowledgments piggybacked on the messages determine a causal order on messages.
- Secure Total Protocol – a total ordering protocol for converting the causal order on messages into a total order ensuring that, even in the presence of crash and Byzantine faults, non-faulty processors determine identical total orders on messages and non-Byzantine processors construct consistent total orders.
- Secure Group Membership Protocol – a membership protocol for maintaining the membership of the configuration by detecting and removing faulty processors and by admitting new and recovered processors.

The Secure Group Membership Protocol operates on top of the Secure Total Protocol, and exploits the total order generated by it. The Secure Total protocol is a probabilistic algorithm in which the probability of making a decision to extend the total order increases asymptotically to unity as more

messages are processed. It is truly fault-tolerant in that it continues to order messages in the presence of both crash and Byzantine faults, provided that a resilience requirement is satisfied. More details about the Secure Group Protocols can be found in Papers 1, 14 and 15.

1.4 The Totem Protocols

The Totem protocols provide reliable totally ordered delivery of messages, as well as topology and membership services, over a local-area network or across multiple local-area networks interconnected by gateways. The Totem protocols employ a logical token-passing ring on each local-area network, with sequence numbers to provide reliable totally ordered delivery of messages on a single local-area network and timestamps to provide total ordering across multiple local-area networks.

The Totem protocols focus on high performance and real time, and undoubtedly have the highest throughput of any ordered multicast group communication protocols yet exhibited, measured for 100 byte messages at 112,000 multicast messages/sec. The Totem protocols allow continued operation in all components of a partitioned network, and ensure that consistency of message ordering is maintained even in the presence of network partitioning faults. The advantage of the Totem protocols over the Trans/Total protocols is that they have a lower computational cost. More details about the Totem protocols can be found in Papers 3, 10 and 12.

1.5 The SecureRing Protocols

The SecureRing Protocols are Byzantine-resistant message ordering and group membership protocols derived from the Totem single-ring protocol.

The SecureRing message ordering protocol protects the system against the following forms of Byzantine attack by a corrupt processor:

- A malicious processor communicates different messages to different destinations, purporting that they are the same message, *i.e.*, giving them the same header. We call these messages *mutants* of each other.
- A malicious processor sends a message which purports to be from another processor.
- A malicious processor alters the token.

To provide protection against these forms of malicious attack, several measures are employed. Each message is digitally signed by its originator to provide authentication. This requirement applies to all messages of any type, including token transmissions. Digests of the previous token and the messages broadcast by the token holder are placed in the token. The token is broadcast, rather than transmitted point-to-point as in the Totem single-ring protocol. These novel strategies allow detection of mutant messages and token alternation.

The SecureRing membership protocol handles benign and malicious processor faults, network partitioning, and loss of the token. When such faults are detected, the membership protocol forms a new ring on which the total ordering protocol can continue operation. To form a new ring, consensus must be reached in that every non-malicious member of the configuration must agree on the membership. Additionally, the membership protocol must terminate, in that every non-malicious processor must install a configuration with an agreed membership within a bounded time unless it fails within that time. To achieve consensus and termination, processors that fail to reach agreement within a bounded time are eliminated from the membership. Byzantine attacks that must be guarded against include:

- A malicious processor causing incorrect membership changes to take place.
- A malicious processor preventing necessary membership changes from taking place.
- A processor that is known to be malicious joining the membership.

In addition to the *proc_set* and *crash_set* used in the Totem single-ring membership protocol, the SecureRing membership protocol employs the *Byz_set*, a set of identifiers of processors that are suspected of being Byzantine.

For a ring of four 200 MHz UltraSPARCs connected over a 100 Mbit/s Ethernet, the measured throughput of the SecureRing Protocols is approximately 900 1000-byte messages/sec, which compares very favorably with the performance of other secure multicast protocols. More details about the SecureRing Protocols can be found in Papers 9 and 20.

1.6 FORTEZZA/Cryptoki vs. Cryptolib

Protocols that can survive Byzantine (arbitrary) faults and malicious attacks come with a high associated overhead. Much of the cost is related to signature generation and verification, which are computationally expensive operations that depend on modular exponentiation.

For the project, we obtained FORTEZZA encryption cards and readers from NSA. We experienced some difficulties initially with the Cryptoki software package from NSA, but these problems were overcome. The FORTEZZA card driven by the Cryptoki software appears to be appropriate up to about 100,000 bit/sec, a significant proportion of the overhead being incurred within the Cryptoki software. Because that performance would substantially limit the performance of our protocols, we abandoned the use of the FORTEZZA cards.

We then obtained from AT&T Labs the Cryptolib software package. Even though it runs entirely in software, Cryptolib runs substantially faster than the hardware FORTEZZA cards on our Sun workstations.

In Table 1 we give some sample execution times using CryptoLib for computing MD4 and MD5 message digests, and for signing and verifying using RSA. These measurements were taken on a 167 MHz Sun UltraSPARC running Solaris 2.5.1. Signatures are computed by RSA decrypting a message digest using the private key, while verification is performed by RSA encrypting the signature using the public key. The signature and verification times shown in Table 1 assume that the message digest has already been computed, and use an RSA key modulus size of 300 bits. Because the message digest is a fixed size (16 bytes), the time required for signing is independent of the size of the original message. However, signature generation time is highly related to key modulus size; a tradeoff exists between performance and the level of security attained.

1.7 The Immune System

The Immune System has emerged from the SecureRing Protocols of this project and from the Eternal System of another DARPA funded project in our Lab. While the SecureRing Protocols handle malicious attack from

Message Size	MD4	MD5	RSA Sign	RSA Verify
250 bytes	30 μs	50 μs	7.4 ms	1.5 ms
500 bytes	50 μs	80 μs		
1000 bytes	100 μs	150 μs		
2000 bytes	200 μs	300 μs		

Table 1: Approximate Times to Compute Digests and Signatures.

within the system, and the Eternal System provides replication management, neither system completely addresses the survivability requirements of critical applications. The Immune System provides intrinsic support for both reliability and security, by the use of object replication, as in the Eternal System, but exploiting instead, the mechanisms of the SecureRing Protocols.

The Immune System can protect an existing unmodified CORBA application, running over an unmodified commercial ORB, against arbitrary faults, including those that arise from malicious attacks within the system. Every object within the CORBA application is actively replicated by the Immune System, with majority voting applied to incoming invocations and responses for each object replica. The Immune System exploits the stringent guarantees of the SecureRing Protocols to enable the majority voting to be effective, even when processors within the network and objects within the application become corrupted.

The Immune System exploits the facilities of the underlying SecureRing Protocols to provide secure reliable totally ordered message delivery. By mapping the intercepted IIOP messages onto the SecureRing Protocols, the Replication Manager of the Eternal System ensures that the client-server interactions are communicated in multicast messages, without modification of either the application objects, or of the ORB.

We have measured the performance of the Immune System for a test application using the VisiBroker 3.2 ORB from Inprise Corporation. The measurements were taken over a network of six dual-processor 167 MHz UltraSPARC workstations, running the Solaris 2.5.1 operating system, connected by a 100 Mbit/sec Ethernet.

For a client generating invocations every 152 μs and for messages ranging in size from 300-1400 bytes (depending on the packing done by the ORB), the measured throughput is as follows. When the secure reliable totally ordered group communication of the underlying protocols is used with message

digests, as well as signatures for the tokens, the measured throughput is 375 messages/sec. When the reliable totally ordered group communication of the underlying protocols is used with message digests alone, the measured throughput is 1840 messages/sec. The cost of digital signatures is a dominant cost in the protocols, even though only the token is signed.

1.8 Achieving Consensus in a Byzantine Environment

Consensus is a fundamental problem in distributed computing that underlies the message ordering and group memberhip algorithms of group communication systems. Fischer, Lynch and Paterson have shown that it is impossible to achieve consensus in an asynchronous distributed system that is subject to even one crash fault. Chandra and Toueg have shown, however, that consensus is possible in an asynchronous distributed system that is subject to crash faults if the asynchronous model is augmented with an unreliable fault detector.

We have investigated the use of fault detectors for solving the consensus problem in asynchronous distributed systems that are subject to Byzantine faults. We capture the essence of Byzantine faults by defining them in terms of deviation from the algorithm A that the processes run. We have defined two new completeness properties, eventual strong Byzantine completeness for algorithm A and eventual weak Byzantine $(k+1)$-completeness for algorithm A, and have used these completeness properties and previously defined accuracy properties to define four new classes of unreliable Byzantine fault detectors.

We have developed an algorithm that uses a Byzantine fault detector to solve the consensus problem in an asynchronous distributed system of n processes subject to at most $\lfloor (n-1)/3 \rfloor$ Byzantine faults. The algorithm employs a rotating coordinator and proceeds in asynchronous rounds. We have also developed an algorithm that implements, in a model of partial sychrony, a fault detector that can be used with our consensus algorithm.

More details on the problem of solving consensus in a Byzantine environment can be found in Papers 9 and 16. The problem of solving membership in a Byzantine environment where processes can fail and recover is considered in Paper 19.

1.9 Real-Time Graphical Interval Logic

Because this project was funded out of the Formal Methods program, we also include here a brief description of our activity in the area of Formal Methods. This work concerns the tools that we have developed for Real-Time Graphical Interval Logic (RTGIL).

The RTGIL tools are intended for specifying and reasoning about time-bounded safety and liveness properties of concurrent real-time systems. These tools include a syntax-directed editor that enables the user to construct graphical formulas on a workstation display, a theorem prover based on a decision procedure that checks the validity of attempted proofs and produces a counterexample if an attempted proof is invalid, and a proof management and database system that tracks proof dependencies and allows graphical formulas to be stored and retrieved. Papers 4 and 5 contain more details on the RTGIL tools.

We have also developed probabilistic duration automata for specifying and analyzing real-time systems in terms of probability density functions. Paper 2 describes this work.

2 Accomplishments vs. Goals

The Secure Group Protocols were implemented by Nitya Narasimhan (item 1.1 on the schedule), including the multicasting of messages, the transitive positive acknowledgment mechanism, the negative acknowledgment and message retransmission mechanisms, the Byzantine resistant voting and total ordering algorithms, and the membership algorithm. However, the protocols are not yet fully tested and are not yet very robust. In particular, no testing for resistance to Byzantine attacks has yet been undertaken, and no performance measurements have been made (Item 2.2). Because the performance of the SecureRing Protocols was substantially better, a decision was made to focus our effort on that protocol.

The SecureRing Protocols were implemented by Kim Kilstrom (Items 1.2 and 2.2 on the schedule). The protocols exploit digital digests to avoid the expense of computing digital signatures for every message, resulting in a substantial performance gain. Messages are now being multicast and ordered efficiently despite the regrettably slow encryption and decryption available to us. Throughput has been measured at 900 1000-byte multicast messages/sec

(which compares with the several seconds per message reported by another researcher at a recent DARPA meeting). The membership algorithm requires considerable care to protect it against Byzantine attacks, but a robust membership algorithm has now been coded and is being tested.

The SecureRing Protocols have been integrated with the Eternal replication manager for CORBA, to produce the Immune System. Eternal provides both active and passive fault replication for CORBA objects using a crash fault model. To extend Eternal with majority voting to protect against commission faults, a commission fault tolerant multicast protocol, such as SecureRing, is necessary (at the same DARPA meeting, another team reported majority-voted fault-tolerance for CORBA, but neglected to protect their multicast protocol against commission faults). The Immune System is part of the basis for a submission to the Object Management Group on Fault Tolerance for CORBA (Item 2.4 on the schedule) in which we have included majority-voted fault-tolerance.

A preliminary version of the library (Item 1.3 on the schedule) now exists, though some of the routines are in C rather than in C++.

The SecureRing Protocols have been demonstrated (Milestone 1.6 on the schedule).

Little work has been done on Items 1.4, 1.5, 2.3, 2.5 or 2.6 on the schedule. Items 1.4 and 1.5 have been set aside for lack of adequate knowledge of, or equipment for, spread spectrum communication in our group. Item 2.3 has been set aside because it is being addressed quite competently by other DARPA researchers investigating intrusion detection. Item 2.5 has not started because the SecureRing Protocols have only recently become operational and there has not been time to adapt it to the Internet or ATM environmants. However, we have two other students, Karlo Berket and Ruppert Koch, working on group communication protocols for the Internet and ATM. We hope to accomplish Item 2.6 over the next nine months as Kim Kihlstrom completes her thesis.

3 Publications

1. Total Ordering Algorithms for Asynchronous Byzantine Systems, L. E. Moser and P. M. Melliar-Smith, *Proceedings of WDAG '95, 9th International Workshop on Distributed Algorithms*, Lecture Notes in Computer Science 972, Le Mont Saint Michel, France, September 1995, 242-256.

2. Probabilistic Duration Automata for Analyzing Real-Time Systems, L. E. Moser and P. M. Melliar-Smith *Proceedings of TACAS'96*, Passau, Germany (March 1996), Lecture Notes in Computer Science 1055, Springer Verlag, 369-390.

3. Totem: A Fault-Tolerant Multicast Group Communication System, L. E. Moser, P. M. Melliar-Smith, D. A. Agarwal, R. K. Budhia, and C. A. Lingley-Papadopoulos, *Communications of the ACM* 39, 4, April 1996, 54-63.

4. The Real-Time Graphical Interval Logic Toolset, L. E. Moser, P. M. Melliar-Smith, Y. S. Ramakrishna, G. Kutty and L. K. Dillon, *Proceedings of the Conference on Computer-Aided Verification*, New Brunswick, NJ, July/August 1996, Lecture Notes in Computer Science 1102, Springer-Verlag, 446-449.

5. A Graphical Environment for Design of Concurrent Real-Time Systems, L. E. Moser, Y. S. Ramakrishna, G. Kutty, P. M. Melliar-Smith and L. K. Dillon, *ACM Transactions on Software Engineering Methodology* 6, 1, January 1997, 31-79.

6. Security in Fault-Tolerant Distributed Systems, K. Kihlstrom, Term Paper, CS 277, Computer Security, University of California, Santa Barbara, Fall 1996.

7. Emerging Technologies in Computer Networks and Distributed Systems, K. Berket, R. K. Budhia, K. P. Kihlstrom, R. Koch, N. Narasimhan, P. Narasimhan, E. M. Royer, M. D. Santos, A. Shum, E. Thomopoulos, P. M. Melliar-Smith and L. E. Moser, *IEEE Looking Forward* 5, 3, September 1997, 2-6.

8. Solving Consensus in a Byzantine Environment Using an Unreliable Fault Detector, K. P. Kihlstrom, L. E. Moser and P. M. Melliar-Smith, *Proceedings of the International Conference on Principles of Distributed Systems*, December 1997, Picardie, France, 61-75.

9. The SecureRing Protocols for Securing Group Communication, K. P. Kihlstrom, L. E. Moser and P. M. Melliar-Smith, *Proceedings of the Hawaii International Conference on System Sciences*, Kona, Hawaii, January 1998, vol. 3, 317-326.

10. The Totem Multiple-Ring Ordering and Topology Maintenance Protocol, D. A. Agarwal, L. E. Moser, P. M. Melliar-Smith and R. K. Budhia, *ACM Transactions on Computer Systems*, vol. 16, no. 2 (May 1998), 93-132.

11. Network Protocols, P. M. Melliar-Smith and L. E. Moser, *Computational Grids: The Future of High-Performance Distributed Computing*, ed. I. Foster and C. Kesselman, Morgan-Kaufmann, 1998.

12. The Totem System, D. A. Agarwal, L. E. Moser and P. M. Melliar-Smith, *Encyclopedia of Distributed Computing*, ed. J. Urban and P. Dasgupta, Kluwer Academic Publishers, 1999.

13. Group Communication, P. M. Melliar-Smith and L. E. Moser, *Encyclopedia of Electrical and Electronics Engineering*, ed. J. Webster, John Wiley & Sons, February 1999.

14. Byzantine-Resistant Total Ordering Algorithms, L. E. Moser and P. M. Melliar-Smith, *Information and Computation*, to appear.

15. The Secure Group Communication System, L. E. Moser, P. M. Melliar-Smith and N. Narasimhan, submitted.

16. Unreliable Byzantine Fault Detectors for Solving Consensus, K. P. Kihlstrom, L. E. Moser and P. M. Melliar-Smith, submitted.

17. Detection of Byzantine Faults in Distributed Systems, K. P. Kihlstrom, L. E. Moser and P. M. Melliar-Smith.

18. Providing Support for Survivable CORBA Applications with the Immune System, P. Narasimhan, K. P. Kihlstrom, L. E. Moser and P. M. Melliar-Smith, submitted.

19. Solving Group Membership with an Unreliable Fault Detector, L. E. Moser, P. M. Melliar-Smith and K. P. Kihlstrom, under revision.

20. Secure Reliable Group Communication with the SecureRing Protocols, K. P. Kihlstrom, L. E. Moser and P. M. Melliar-Smith, in preparation.

21. Fault Tolerance for CORBA, Version 1.0, Initial RFP Submission, OMG Document orbos/98-10-08, L. E. Moser, P. M. Melliar-Smith and P. Narasimhan, Technical Report 98-27, Department of Electrical and Computer Engineering, University of California, Santa Barbara (October 1998).

4 Professional Personnel

4.1 Professors

Louise Moser
Michael Melliar-Smith

4.2 Ph.D. Students

Kim Kihlstrom
Nitya Narasimhan

5 New Discoveries, Inventions and Patents

The new discoveries are described in Section 1 and in the papers appended to this report. There were no inventions on the project that the investigators consider patentable.

6 Other Activities

6.1 Meetings Attended

L. E. Moser and P. M. Melliar-Smith, DARPA Survivability-Formal Methods PI meeting, San Diego, January 1996

P. M. Melliar-Smith, ARPATECH'96, Atlanta GA, May 1996

L. E. Moser and P. M. Melliar-Smith, DARPA PI Meetings, Dallas, TX October 1996 and December 1996

L. E. Moser and P. M. Melliar-Smith, DARPA PI Meeting, Washington, DC, July 1997

L. E. Moser and P. M. Melliar-Smith, DARPA PI Meeting and Workshop on Composition and Wrappers, Lake Tahoe, CA, August 1997

L. E. Moser and P. M. Melliar-Smith, DARPA-OMG-MCC Workshop on Compositional Software Architectures, Monterey, CA, January 1998

L. E. Moser and P. M. Melliar-Smith, DARPA PI Meeting, Menlo Park, CA, May 1998

L. E. Moser, P. M. Melliar-Smith, R. Koch and M. Santos, DARPA PI Meeting, San Diego, CA, July 1998

6.2 Presentations

L. E. Moser, Secure Multicast Protocols for Group Communication, DARPA Survivability-Formal Methods PI meeting, San Diego, CA, January 1996

L. E. Moser, Probabilistic Duration Automata for Analyzing Real-Time Systems, 2nd International Workshop on Tools and Algorithms for the Construction and Analysis of Systems, Passau, Germany, March 1996

L. E. Moser, P. M. Melliar-Smith and students, Fault-Tolerant Group Communication for Clusters of Computers, Presentation and Software Demonstration, Sun Microsystems, Menlo Park, CA, June 1996

L. E. Moser and P. M. Melliar-Smith, Secure Multicast Protocols for Group Communication, DARPA Formal Methods PI meeting, Lake Placid, NY, July 1996

L. E. Moser, P. M. Melliar-Smith and Y. S. Ramakrishna, The Real-Time Graphical Interval Logic Toolset, Visual Reasoning Workshop, Conference on Automated Deduction, New Brunswick, NJ, July 1996

L. E. Moser, P. M. Melliar-Smith and Y. S. Ramakrishna, The Real-Time Graphical Interval Logic Toolset, Computer Aided Verification Conference, New Brunswick, NJ, August 1996

L. E. Moser, Group Communication for Fault-Tolerant Distributed Systems, Florida State University, Tallahassee, FL, January 1997

L. E. Moser, Group Communication for Distributed Networked Systems, San Jose State University, San Jose, CA, March 1997

L. E. Moser and P. M. Melliar-Smith, The Totem System, Demonstration, DoD Research Advocacy Day, DARPA, Washington, D.C., May 1997

L. E. Moser, Protecting Distributed Systems against Byzantine Attacks, DARPA PI Meeting and Workshop on Composition and Wrappers, Lake Tahoe, CA, August 1997

L. E. Moser, The Eternal System, DARPA Workshop on Fault Tolerance, JPL, Pasadena, CA, September 1997

L. E. Moser, P. M. Melliar-Smith, K. Kihlstrom, P. Narasimhan, N. Narasimhan, V. Kalogeraki, L. Tewksbury, Fault-Tolerant Distributed Systems, Jet Propulsion Laboratory, November 1997.

K. P. Kihlstrom, Solving Consensus in a Byzantine Environment Using an Unreliable Fault Detector, International Conference on Principles of Distributed Systems, Picardie, France, December 1997.

K. P. Kihlstrom, The SecureRing Protocols for Securing Group Communication, Hawaii International Conference on System Sciences, Kona, Hawaii, January 1998.

L. E. Moser and P. M. Melliar-Smith, The Eternal and Realize Systems, DARPA-OMG-MCC Workshop on Compositional Software Architectures, Monterey, CA, January 1998

L. E. Moser, P. M. Melliar-Smith and R. Koch, Strategies for Building Fault-Tolerant Distributed Systems, Presentation and Software Demonstration, Channel Islands Chapter, American Society of Naval Engineers, Pt Mugu, February 1998

L. E. Moser, The Eternal System, DARPA PI Meeting, Menlo Park, CA, May 1998

L. E. Moser, P. M. Melliar-Smith, R. Koch and M. Santos, Software Demonstration, The Realize System, Totem System and Atomic Group System, DARPA PI Meeting, San Diego, CA, July 1998

L. E. Moser, P. M. Melliar-Smith and their students made presentations and gave software demonstrations for the following visitors to the project.

6.3 Visitors to the Project

Professor Klaus Petermann, Technical University of Berlin

Professor Ben Wah, University of Illinois

Dr. Kevin C. Almeroth, Georgia Institute of Technology

Dr. Ender Ayanoglu, Lucent/Bell Labs

Dr. Rachid Guerraoui, Ecole Polytechnique Federale de Lausanne

Dr. David Blumenthal, Georgia Institute of Technology

Richard Thibault and Bruce Canna, The Foxboro Company

Professor Dan Gajski, University of California, Irvine

Professor Douglas Schmidt, Washington University, St. Louis

Dr. Michael Reiter, AT&T Laboratories

Dr. Gil Neiger, Intel

Dr. Brian A. Hanson, Director, Speech Technology Laboratory Panasonic Technologies, Inc, with 12 other Directors of Panasonic's USA Laboratories

Keith Bromley, NRad, San Diego

Dennis Hollingworth, Trusted Information Systems

Dr. C. K. Toh, Hughes Research Laboratory

Brian Norling, Director of Engineering, Space and Launch Systems, Litton Guidance and Control

Dr. Hossein Moiin, Sun Computer Company

Dr. Deborah Agarwal, Lawrence Berkeley National Laboratory

Professor Partha Dasgupta, Arizona State University

Professor Hermann Kopetz, Technical University of Vienna, Austria

James Kirkley, John Norris, Brian Whittle and Chad Stone, QAD, Carpenteria, CA

David Lomet, Microsoft Corporation

Dr. Lewis B. Oberlander, Dr. Won Kang, Francis Tam, Motorola Corporation, Celluar Infrastructure Group, Arlington Heights, IL

Dr. Gregory Papodopolous (Chief Technology Officer), Dr. Emil Sarpa (Director of External Research), Dr. John M. Hale (Program Manager External

Research) and Georgi C. Johnson (Technology Collaboration Manager), Sun Microsystems

Dr. Stephen Wright, Mathematics and Computer Science Division, Argonne National Laboratory

Professor Georgis Giannakis, University of Virginia

Greg Hoffman and Mihir Ravel, Tektronix

Mark Gibbs and John Dix, Network World

Mike Toma and Vic Walker, jeTech Data Systems, Inc, Camarillo, CA

Professor Odd Pettersen, Norwegian University of Science and Technology

6.4 Industrial Interest

Discussions have been held with Deborah Agarwal of Lawrence Berkeley National Laboratory on the multicast delivery services needed for DOE's Collaboratories, which will enable scientists to collaborate over the Internet. We currently have a project with LBNL to implement multicast protocols for the Collaboratories. Our student, Nitya Narasimhan, worked at LBNL during the summer of 1997.

At the DARPA meeting in Dallas, Texas, in December 1996, discussions were held with Terry Benzel of Trusted Information Systems. Dennis Hollingworth from Trusted Information Systems visited our Lab in November 1997. Dennis found the Secure Multicast Protocols fascinating and, as a result, an article on the SecureRing Protocols appeared in the TIS Newsletter.

Information on the SecureRing and Secure Group Protocols, as well as on our experience with the FORTEZZA cards and Cryptolib software, has been provided to Rob Ruth.

Keith Bromley of NRad visited our Lab in October 1997. He had heard Louise Moser's presentation, Protecting Distributed Systems against Byzantine Attack, at the DARPA meeting in Washington, D.C. in July 1997, and had recommended that we present our work at the Jet Propulsion Laboratory to the researchers involved in the next generation space program X2000. Our work on Byzantine (arbitrary) faults is of interest to the JPL researchers

because of the arbitrary faults they experience in their computer systems due to the bombardment by alpha particles in space.

Louise Moser made a presentation in September 1997 at the DARPA sponsored workshop on fault tolerance at JPL. Louise Moser and Michael Melliar-Smith and five of their students made a second presentation at JPL in November 1997. JPL is very interested in using the technology being developed by Moser and Melliar-Smith and their students, both in this project and in our other DARPA sponsored projects.

In December 1997 we presented recent work at Sun Microsystems in Menlo Park, and Greg Papodopolous, the Chief Technical Officer of Sun Microsystems, and some of his colleagues visited our laboratory in March 1998. In his subsequent debriefing with the Dean of Engineering at UCSB, Dr. Papodopolous noted the remarkably close correspondence between the research being undertaken by us and the research needs of Sun Microsystems. We expect collaborations with Sun Microsystems to continue.

The RFP issued by the Object Management Group for Fault Tolerance in CORBA contains a requirement for fault tolerance by majority voting, so that commission faults can be masked in addition to crash faults. Masking commission faults requires not only majority voting to mask commission faults in the application objects but also a reliable multicast protocol to mask commission faults that affect the multicast protocol, as in the Immune System. Our proposal, submitted to OMG in response to their RFP on Fault Tolerance in CORBA, includes majority voting and also ân explanation of the need for multicast protocols that can resist commission faults, such as the SecureRing protocol. A copy of the proposal is included in this report.

MISSION
OF
AFRL/INFORMATION DIRECTORATE (IF)

The advancement and application of Information Systems Science and Technology to meet Air Force unique requirements for Information Dominance and its transition to aerospace systems to meet Air Force needs.